ABSTRACT OF THE DISCLOSURE

A load balancing technique uniformly distributes data transmitted by a server over a plurality of underlying physical links of a virtual interface or aggregate within a computer network. According to the invention, the unit of data granularity is preferably a datagram, rather than an Internet protocol (IP) fragment or packet. The load balancing technique implements a round robin scheduling policy among the links based on an identifier (ID) associated with each datagram. As a result, the technique ensures that all fragments of a datagram are forwarded over the same physical link of the aggregate on the basis of a similar IP ID.